

# Huge-scale Molecular Dynamics Simulation of Multibubble Nuclei

Hiroshi Watanabe<sup>\*1</sup>, Masaru Suzuki<sup>2</sup>, and Nobuyasu Ito<sup>3</sup>

<sup>1</sup>*The Institute for Solid State Physics, The University of Tokyo,  
Kashiwanoha 5-1-5, Kashiwa, Chiba 277-8581, Japan*

<sup>2</sup>*Department of Applied Quantum Physics and Nuclear Engineering,  
Kyushu University, 744 Motooka, Nishi-ku, Fukuoka 819-0395, Japan and*

<sup>3</sup>*Department of Applied Physics, School of Engineering,  
The University of Tokyo, Hongo, Bunkyo-ku, Tokyo 113-8656, Japan*

We have developed molecular dynamics codes for a short-range interaction potential that adopt both the flat-MPI and MPI/OpenMP hybrid parallelizations on the basis of a full domain decomposition strategy. Benchmark simulations involving up to 38.4 billion Lennard-Jones particles were performed on PRIMEHPC FX10, consisting of 4800 SPARC64 IXfx 1.848 GHz processors, at the Information Technology Center of the University of Tokyo, and a performance of 193 teraflops was achieved, which corresponds to a 17.0% execution efficiency. Cavitation processes were also simulated on PRIMEHPC FX10 and SGI Altix ICE 8400EX at the Institute of Solid State Physics of the University of Tokyo, which involved 1.45 billion and 22.9 million particles, respectively. Ostwald-like ripening was observed after the multibubble nuclei. Our results demonstrate that direct simulations of multiscale phenomena involving phase transitions from the atomic scale are possible and that the molecular dynamics method is a promising method that can be applied to petascale computers.

Keywords: Hybrid MPI + OpenMP programming; Molecular Dynamics Method;

## I. INTRODUCTION

Most recent computers consist of multiple nodes, with each node usually having multiple processors. Additionally, not only state-of-the-art computer but commodity cluster systems adopt hierarchical memory models, *i.e.*, shared memory within a node and distributed memory across the nodes. Corresponding to the different memory models, there are two programming models, the Message Passing Interface (MPI) library and OpenMP, which are the de facto standards for threading distributed and shared memory models, respectively. To achieve better performance on systems with such a heterogeneous memory model, programmers have to choose their strategies for parallelization from the following three possible approaches; fully distributed memory models, global address space models, and a combination of distributed and shared memory models. The fully distributed memory models are achieved by the so-called flat-MPI strategy, *i.e.*, shared memory in a node is treated as distributed memory and only the MPI library is used for both intra- and internode communication. This strategy reduces development costs, since a programmer can ignore the heterogeneity of the memory model in the system. Also, the flat-MPI approach usually achieves acceptable performance since the MPI library is usually optimized so that it can utilize a shared memory model within a node. The global address space models allow programmers to use a virtually global memory address space, which is distributed over nodes. Using partitioned global address space (PGAS) languages, programmers can control the affinity of the memory space to threads. There are many languages that adopt this programming model, for example, Unified Parallel C [1], Co-Array Fortran [2], Titanium [3], XcalableMP [4], and so forth. While these languages reduce the development cost of software, users have to utilize their features to the full to obtain sufficient performance for large-scale computations. Finally, the combination of distributed and shared memory models is usually achieved by using the MPI library for intranode communication and OpenMP for intranode communication. Although it seems natural to adopt the distributed memory model across nodes and the shared memory model within each node, the flat-MPI approach usually gives better performance than the hybrid approach. However, one cannot always adopt the flat-MPI approach since it requires more resources such as the total memory, especially for massively parallel simulations. Many comparisons between the flat-MPI approach and the OpenMP/MPI hybrid approach have been carried out. However, OpenMP/MPI hybrid programming is usually implemented by domain decomposition with the MPI library and loop-level decomposition with OpenMP. Even if the total systems of two schemes are identical, the computations applied to each core are different. This makes difficult to compare the efficiency between the two schemes, since it is difficult to find reasons for the difference in performance.

In this paper, we describe a pseudo-flat-MPI implementation of molecular dynamics (MD) simulations that allows us to compare the performance between flat-MPI and a hybrid using the same program. We adopt the domain

---

\* hwatanabe@issp.u-tokyo.ac.jp

decomposition strategy for both intra- and internode parallelizations. We also focus on how to perform product runs using the hybrid parallelized MD, since the development cost of implementing codes for product runs can be compared with that of implementing benchmark codes. Programmers have to consider many factors, such as the parallelization of observation codes, file IOs, and memory consumption. The scaling analysis of benchmark simulations involving up to 38.4 billion Lennard-Jones (LJ) particles is performed on PRIMEHPC FX10, consisting of 4800 SPARC64 IXfx 1.848 GHz processors, at the Information Technology Center of the University of Tokyo. To achieve better performance on FX10, we optimize the simulation kernel and achieve a performance of 193 teraflops, which corresponds to a 17.0% execution efficiency. We also perform substantial simulations of cavitation phenomena involving up to 1.45 billion particles on two different platforms, FX10 and SGI Altix ICE 8400 EX (Intel Xeon 2.93 GHz processors), at the Institute of Solid State Physics of the University of Tokyo.

This paper is organized as follows. In Sec. II, details of our method are described including the parallelization. Benchmark results are given in Sec. III and cavitation phenomena are studied in Sec. IV. Finally, Sec. V is devoted to a summary and discussions of this study.

## II. METHOD

### A. Basic Algorithms

We perform classical MD simulations of LJ potentials with truncation. Here, we briefly describe the algorithms for short-range MD. Refer to a previous paper for detailed descriptions [5]. The two-body potential of truncated LJ particles is given by [6]

$$V(r) = 4\epsilon \left[ \left( \frac{\sigma}{r} \right)^{12} - \left( \frac{\sigma}{r} \right)^6 + c_2 \left( \frac{r}{\sigma} \right)^2 + c_0 \right], \quad (1)$$

with particle distance  $r$ , well depth  $\epsilon$ , and atomic diameter  $\sigma$ . The two additional coefficients,  $c_2$  and  $c_0$ , are determined so that the potential and force become zero at the cutoff length  $r_c$ , *i.e.*,  $V(r_c) = V'(r_c) = 0$ . In the following, we use physical quantities reduced by  $\sigma$ ,  $\epsilon$ , and the Boltzmann constant  $k_B$ , *i.e.*, the length scale is measured using the unit of  $\sigma$ , and so forth. We set the cutoff length to  $r_c = 2.5$  for benchmark simulations and to  $r_c = 3.0$  for simulations of cavitation. Since interactions between particles are short-range, it is necessary to find particle pairs separated by a distance of them is less than the cutoff length. We adopt a grid algorithm [7–10] (also called the linked-list method [11, 12]) to find the interacting particle pairs. Since the computational cost of constructing a list of interacting particle pairs is high, we adopt the bookkeeping method to reduce the cost of constructing the list [13], which allows us to reuse the same pair list for several time steps by registering pairs within some length which is longer than the cutoff length. Several steps after the construction of the pair list, there may be particle pairs separated by a distance less than the cutoff length that are not registered in the list. If such pairs exist, the simulation fails since the total energy will not be conserved. Therefore, we have to check the validity of the list in every time step. For this, we adopt the dynamic upper time cutoff (DUTC) method [14]. The validity of the list is checked using the maximum velocity and the displacements of the particles [5]. This validity check involves the global synchronization between all processes.

After constructing a pair list, we sort the list to reduce the number of accesses to the memory (see Fig. 1). Consider a particle pair  $(i, j)$ , where  $i < j$ . We call the particle of index  $i$  the *key* particle and the other the *partner* particle. We sort the particle pairs by the indices of the *key* particles. Then the data of the *key* particle is stored in CPU registers, and the amount of memory access is decreased since data fetching and storing are performed only for the data of *partner* particles.

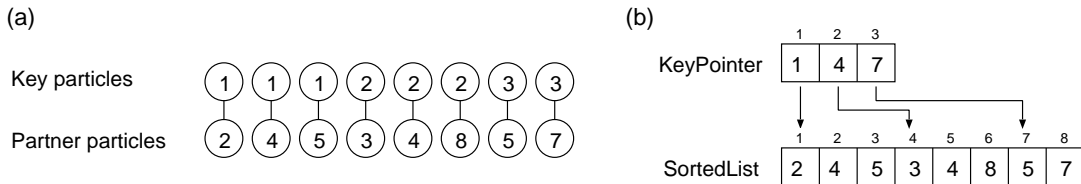


FIG. 1: (a) Sorted pair list. The pair list is sorted by the indices of *key* particles. (b) Data format of the sorted pair list. A sorted pair list is expressed by two arrays; KeyPointer and SortedList. The first partner of *key* particle  $i$  is stored in  $\text{SortedList}[\text{KeyPointer}[i]]$ , the second is stored in  $\text{SortedList}[\text{KeyPointer}[i]+1]$ , and so forth.

## B. Parallelization

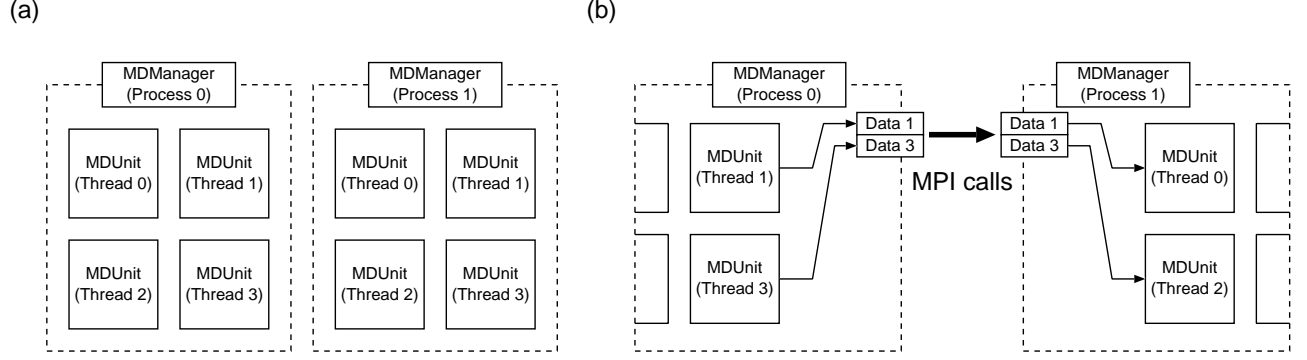


FIG. 2: Schematic illustration of the pseudo flat-MPI approach. A two-dimensional system is shown for convenience. There are two MPI processes, each containing four OpenMP threads. (a) MDUnit and MManager represent an MPI process and an OpenMP thread, respectively. Each domain is assigned to an MDUnit, and an MManager takes responsibility for their communication. (b) Communication between processes. Since a platform may not support MPI calls from a multithread environment, data should be packed before being sent by a sender process, and the data should be distributed to destinations by a receiver process.

Although there are several strategies for the parallelization of MD [15], the domain decomposition method is simple and generally scalable for a huge number of processes. It is also the most efficient for particle systems with short-range interactions if the interaction length is sufficiently shorter than the size of the domain assigned to each process. Domain decomposition parallelization is usually implemented by the flat-MPI approach. This approach allows a programmer to consider only a single programming model, and therefore, the development cost is reasonable for huge scale computations. However, limited resources, usually a lack of memory, may prevent the adoption of the flat-MPI programming model. Thus, the hybrid programming model is required to reduce the number of MPI processes. One method of implementing hybrid parallelized MD is to adopt domain decomposition for internode communication with the MPI and particle decomposition for intranode communication with loop-level OpenMP. One advantage of this method is that the implementation can be simple. By inserting an OpenMP directive before the loop of force calculations, one can modify a flat-MPI code to an OpenMP/MPI hybrid code. However, there are some disadvantages of this method. Construction of the pair list typically costs 10% of the total computation time. If the parallelization efficiency of the force computation within a node is perfect, then the construction of the pair list may become a new bottleneck in simulations. Therefore, one has to implement OpenMP parallelization of the pair list construction, as well as the force calculation. Additionally, the observation can also be a bottleneck. Therefore, a programmer has to consider all routines and implement hybrid versions of them if necessary. This pushes up the development cost.

In this study, we adopt a pseudo-flat-MPI approach, *i.e.*, we adopt a full domain decomposition strategy for both intra- and internode parallelizations by using thread IDs as ranks of the MPI [16]. The advantages of the pseudo-flat-MPI approach as follows, with details will be described later.

**Simple implementation:** If one already has a flat MPI code, it is straightforward to modify it to a pseudo-flat-MPI code.

**Flexibility:** It is easy to change the number of threads per process without rebuilding the program.

**Affinity with CPU-level tuning:** For the naive loop-level OpenMP paradigm, one has to consider OpenMP parallelization and SIMD (single instruction multiple data) optimization simultaneously. The pseudo-flat-MPI paradigm allows a programmer to consider only SIMD optimization.

**Memory locality:** For systems with Non-Uniform Memory Access (NUMA) architecture, a programmer has to consider memory locality, *i.e.*, which data should be assigned to which CPU core. In the pseudo-flat-MPI paradigm, the memory locality is automatically satisfied as for the flat-MPI paradigm.

In the following, we describe our implementations in C++ language since we have written our codes in C++. However, similar approaches are possible in C or Fortran 77 or other non object-oriented languages. We define two classes; MManager and MDUnit. An MDUnit class represents an OpenMP thread (see Fig. 2 (a)). The system is

divided into small domains, and one domain is assigned to one instance of the `MDUnit` class, *i.e.*, the instance of the `MDUnit` class is the owner of the domain and is responsible for updating the momenta and coordinates of particles in the domain. An `MDManager` class represents an MPI process. It manages at least one instance of `MDUnit` and is responsible for communication between domains. Sample codes for the initialization of `MDUnit` are as follows.

```
MDUnit *mdp;
std::vector<MDUnit*> mdv;
#pragma omp parallel shared(mdv) private(tid,mdp)
{
    tid = omp_get_thread_num();
    mdp = new MDUnit(tid);
#pragma omp critical
    mdv.push_back(mdp);
}
```

Since variables are initialized in each constructor of `MDUnit`, the memory locality is automatically satisfied. The number of threads per process corresponds to the number of instances of `MDUnit` managed by an instance of `MDManager`. If one `MDManager` has only one instance of `MDUnit`, then the computation is performed in the manner of the flat-MPI paradigm; otherwise it is performed in the manner of the hybrid. The number of threads per process is easily changed by changing the value of the `OMP_NUM_THREADS` environment variable. A sample code for calculating force is as follows.

```
#pragma omp parallel for schedule(static)
for(int i = 0; i < num_threads; i++) {
    mdv[i]->CalculateForce();
}
```

Here, `num_threads` is the number of threads per process. The implementation of the `MDUnit` class does not depend on whether it is called by a thread or a process except is the case of communication. Therefore, it is easy to modify an existing flat-MPI code to a pseudo-flat-MPI code.

Communication involving hybrid parallelization should be considered since there are several levels of support for a call of MPI routines in a multithread environment. The specifications of MPI 2.2 define the following four levels of support [17].

- **MPI\_THREAD\_SINGLE**: MPI calls in the OpenMP parallel region are not supported.
- **MPI\_THREAD\_FUNNELED**: The process may be multithread, but the application must guarantee that only the main thread makes MPI calls.
- **MPI\_THREAD\_SERIALIZED**: The process may be multithread and any thread may make MPI calls. However, the application must guarantee only one call at a time, *i.e.*, consecutive MPI calls from different threads should be serialized.
- **MPI\_THREAD\_MULTIPLE**: Any thread may make MPI calls without any restrictions.

If the system supports `MPI_THREAD_MULTIPLE`, the implementation of communication in the pseudo-flat-MPI becomes identical to that in the flat-MPI. Each instance of `MDUnit` can communicate with others as if it is a process while it is a thread. However, most platforms do not support `MPI_THREAD_MULTIPLE`. For example, the support level of both sites we use in this study is `MPI_THREAD_SERIALIZED`. Therefore, a programmer has to explicitly manage intra- and internode data exchanges. The data required from other processes should be packed before being sent by a sender process and they are distributed by a receiver process (see Fig. 2 (b)). Note that one does not have to pack the data for communication if the system supports `MPI_THREAD_SERIALIZED`. However, making MPI calls outside OpenMP parallel regions is safe and portable, since it will work on platforms that support only `MPI_THREAD_SINGLE`. This two-level communication is also required for observation routines. Therefore, one has to rewrite all codes for observations even if one has already developed parallelized observation codes for a flat-MPI program. This is one of the disadvantages of the pseudo-flat-MPI approach.

We do not consider any load balancing. The dimensions of each domain assigned to a thread are identical among domains and do not change throughout the simulations. Although load imbalance may decrease the efficiency of simulations, the decrease is not serious for our simulations. The reason for this is described in Sec. IV D.

Facility	System	CPU	Cores	Memory	Nodes
ITC	PRIMEHPC FX10	SPARC64 IXfx (1.848 GHz)	16	32 GB	4800
ISSP	SGI Altix ICE 8400EX	Intel Xeon X5570 (2.93 GHz)	8	24 GB	1920

TABLE I: Summary of (ITC) PRIMEHPC FX10 at the Information Technology Center of the University of Tokyo and (ISSP) SGI Altix ICE 8400EX at the Institute for Solid State Physics of the University of Tokyo. The facility, name, CPU, number of cores per node, memory per node, and number of nodes in the system are shown for each system. Although the total number of nodes at ISSP is 1920, we only use 128 nodes in this study.

### III. BENCHMARK RESULTS

#### A. Details of Platforms

We perform simulations on two platforms. One is PRIMEHPC FX10 at the Information Technology Center (ITC), the University of Tokyo. Each node contains a SPARC 64 IXfx 1.848 GHz processor containing 16 CPU cores. The memory access is uniform from each core, *i.e.*, this system adopts Uniform Memory Access (UMA). Each node contains 32 GB and the total amount of memory is 150 TB. We used Fujitsu C/C++ Compiler Ver. 1.2.1 with the compile options `-Kfast -Ksimd=2 -Knoparallel -Kopenmp`. Note that the compile options supported by the Fujitsu C/C++ compiler can be different for different sites.

The other is SGI Altix ICE 8400EX at the Institute for Solid State Physics (ISSP), of the University of Tokyo. Each node contains two Intel Xeon 2.93 GHz processors containing 4 CPU cores, and each node adopts cache-coherent NUMA (ccNUMA). The total system contains 1920 nodes and provides 180 teraflops. Each node contains 24 GB and the total amount of memory is 46 TB. We used Intel C++ Compiler Ver. 11.1 with the compile options `-O3 -ip -ipo -xSSE4.2 -axSSE4.2 -openmp`. A summary of the platforms is shown in Table I.

Although we perform product runs of cavitation phenomena on both FX10 and SGI Altix ICE 8400EX, we perform benchmark simulations only on FX10 at ITC, since the benchmark results on ISSP have already been reported [5].

#### B. Comparison between Flat MPI and Hybrid

The conditions for our benchmark simulations are as follows [5]. All quantities are measured in LJ parameters, such as the particle diameter  $\sigma$ , and the energy scale  $\epsilon$ .

- The dimensions of each domain are  $100 \times 100 \times 100$  and one domain is assigned to one MPI process (flat MPI) or one OpenMP thread (hybrid).
- The number density is 0.5, *i.e.*, each domain contains 500,000 particles.
- Initial condition: Face-centered-cubic lattice.
- Boundary condition: Periodic for all axes.
- Integration scheme: Second-order symplectic integration.
- Time step: 0.001.
- Cutoff length: 2.5.
- Cutoff scheme: Add constant and quadratic terms to potential as shown in Eq. (1).
- Initial velocity: The absolute value of the velocity of all particles is set to 0.9 and the directions of the velocities are given randomly.
- After 150 steps, we measure the calculation (elapsed) time for the next 1000 steps.

Computational speeds are measured in the unit MUPS, which is millions of updates per second. When a system with  $N$  particles is simulated for  $k$  steps in  $t$  seconds, then the number of MUPS is given by  $10^{-6}Nk/t$ . The number of nodes is increased while keeping the same number of particles on each node (weak scaling). We apply two paradigms, flat-MPI and the OpenMP/MPI hybrid for each size of simulation. In the flat-MPI case, we assign 16 processes on a node (one MPI process to one CPU core). In the hybrid case, we assign one MPI process containing 16 OpenMP

threads on a node (one OpenMP thread to one CPU core). We use identical executable file for the two models and we assign domains to nodes so that the same part of the system is assigned to the same CPU core for the flat-MPI and hybrid. Therefore, the simulations are completely identical for the flat-MPI and hybrid runs including the communication patterns across nodes.

The results of the flat-MPI are summarized in Table II and those of the hybrid are summarized in Table III. The elapsed times of the runs are shown in Fig. 3. Since we perform weak scaling analysis, the elapsed time will be independent of the number of nodes for perfect parallel efficiency. As shown in Fig. 3, while the parallel efficiency of flat-MPI is almost perfect, that of the hybrid runs fluctuates. The performance of flat-MPI is always better than that of the hybrid. Since the computations are identical for the two methods, the difference originates from the parallel overheads of OpenMP. There are several possible causes of the overheads. One of them is the cost of thread management. The creation and destruction of threads are required for OpenMP. Additionally, the number of synchronizations per step is larger than in flat-MPI. Although only one global synchronization per step is required for a flat-MPI, several intranode synchronizations are required for the hybrid. Even for a perfect load balance, frequent synchronization may cause parallel overhead due to noise from the system. Note that the total memory required for the hybrid is less than that for flat-MPI, as expected. For the largest run involving 4800 nodes, the hybrid run consumed 10 GB per node out of 32 GB while flat-MPI consumed 14 GB.

Nodes	Number of Particles	Elapsed Time [s]	Speed [MUPS]	Efficiency
1	8,000,000	445.815	17.9447	1.0
72	576,000,000	448.275	1284.93	0.99
480	3,840,000,000	452.781	8480.93	0.98
1440	11,520,000,000	456.641	25227.7	0.98
4800	38,400,000,000	458.399	83769.9	0.97

TABLE II: Results of benchmark simulations of flat-MPI. The number of nodes, number of particles, elapsed time [s], computational speed [MUPS], and parallel efficiency are listed. The parallel efficiency is determined relative to the elapsed time for the single-node calculation.

Nodes	Number of Particles	Elapsed Time [s]	Speed [MUPS]	Efficiency
1	8,000,000	585.160	13.6715	1.00
72	576000000	627.449	918.003	0.93
480	3840000000	724.139	5209.34	0.79
1440	11520000000	724.139	15908.5	0.81
4800	38,400,000,000	684.871	56068.9	0.85

TABLE III: Same as Table II for hybrid runs.

### C. Optimization for FX10

The results shown in the previous section are obtained using a code optimized for Intel Xeon architectures. The performance of the largest job is 56.5 teraflops, which corresponds to an approximately 5% execution efficiency. To improve the efficiency and performance of simulations, we optimize the force calculation routine for FX10. The algorithm optimized for FX10 is given as Algorithm 1, where  $N$  is the total number of particles,  $\mathbf{q}[i]$  is the position of particle  $i$ ,  $\mathbf{p}[i]$  is the momentum, and  $dt$  is the time step. The variables in bold denote three-dimensional vectors, *e.g.*,  $\mathbf{q} = (q_x, q_y, q_z)$ . The constants  $c_2$  and  $c_0$  are coefficients introduced for the truncation in Eq. (1). We find that store operations involving indirect addressing is expensive for FX10. Therefore, we avoid random storing by not using Newton's third law. While the number of iterations required to calculate forces (the length of SortedList) is doubled without Newton's third law, the computational efficiency increases more than twofold. We also utilize SIMD operations. The variables with hats in Algorithm 1 denote the packed variables. A packed variable is stored in 128-bit registers and contains two double-precision 64-bit numbers. Suppose there are two double-precision numbers  $a$  and  $b$  and a packed variable  $\hat{c}$ . Packing  $a$  and  $b$  into  $\hat{c}$  is denoted as  $\hat{c} \leftarrow \{a, b\}$  and unpacking  $\hat{c}$  to  $a$  and  $b$  is denoted as  $\{a, b\} \leftarrow \hat{c}$ . To hide the latency by software pipelining, the inner loop is unrolled four times. Since the two-body potential is truncated at the cutoff length  $r_c$ , the force calculation involves checking whether or not the distance between a particle pair is shorter than the cutoff length. To perform the check with SIMD operations, a new variable,

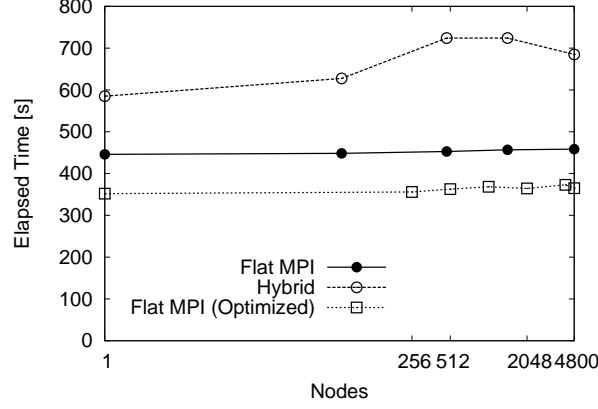


FIG. 3: Results of benchmark simulations for flat-MPI and hybrid runs (weak scaling). Elapsed times for 1000 steps are shown. Although the elapsed times of flat-MPI are almost independent of the number of nodes, the results of hybrid runs fluctuate. Additionally, computational speed of the flat-MPI is always higher than that of the hybrid. The data of flat-MPI using the optimized code for FX10 is also shown. The optimized code is about 25% faster than the original code.

$w = 1 - \lfloor r^2/r_c^2 \rfloor$ , is introduced, where  $\lfloor x \rfloor$  denotes the largest integer not larger than  $x$ . The variable  $w$  can be used as a mask for the truncation of the potential, since  $w = 1$  when the distance between the particle pair is shorter than the cutoff length and  $w = 0$  otherwise. Note that each distance between a particle pair in a pair list should be less than  $\sqrt{2}r_c$  to use this mask. We construct a pair list so that this condition is always satisfied.

The force calculation of the LJ interaction involves at least one division for each particle pair. However, the floating-point divide instruction (fdivd) in FX10 cannot be specified as a SIMD instruction. Therefore, we use the floating-point reciprocal approximation instruction (frcpad) instead of the divide instruction. The reciprocal approximation instruction is faster than the divide instruction at the expense of precision. The rounding error due to the approximation is up to  $1/256$ , *i.e.*,

$$\left| \frac{\text{frcpad}(x) - 1/x}{1/x} \right| < \frac{1}{256}, \quad (2)$$

which is insufficient for MD simulations. Therefore, we improve the precision by the following error correction. Suppose  $r$  is the particle distance. The force calculation of 6-12 LJ interactions requires the values of  $r^{-8}$  and  $r^{-14}$ . Consider the calculation of  $r^{-8}$  from  $r^2$ . Let  $\tilde{r}^{-2}$  be a reciprocal approximation of  $r^2$ , *i.e.*,  $\tilde{r}^{-2} = \text{frcpad}(r^2)$ . The rounding error due to the approximation can be expressed as

$$\tilde{r}^{-2}r^2 = 1 + \varepsilon, \quad (|\varepsilon| < 1/256) \quad (3)$$

with error  $\varepsilon$ . Then the required value  $r^{-8}$  is expressed as

$$r^{-8} = \frac{(\tilde{r}^{-2})^4}{(1 + \varepsilon)^4} \quad (4)$$

$$= (\tilde{r}^{-2})^4(1 - 4\varepsilon + O(\varepsilon^2)) \quad (5)$$

$$= (\tilde{r}^{-2})^4(5 - 4(1 + \varepsilon) + O(\varepsilon^2)) \quad (6)$$

$$= (\tilde{r}^{-2})^4(5 - 4\tilde{r}^{-2}r^2 + O(\varepsilon^2)). \quad (7)$$

Therefore, the error of the corrected value  $(\tilde{r}^{-2})^4(5 - 4\tilde{r}^{-2}r^2)$  relative to  $r^{-8}$  is  $O(\varepsilon^2)$ . Similarly, the correction coefficient for  $r^{-14}$  is obtained to be  $(8 - 7\tilde{r}^{-2}r^2)$ . The effect of the error correction is shown in Fig. 4. One can see that the conservation of the total energy is improved by the error correction.

The performance of the optimized code is summarized in Table IV and shown in Fig. 3 as the data denoted by Flat MPI (Optimized). The increase in speed due to optimization is about 25% compared with the original code with flat MPI. Although the performance also increases to 193 teraflops, which corresponds to a 17.0% execution efficiency, one has to note that it contains redundant computations owing to the nonuse of Newton's third law.

**Algorithm 1** Code for calculating force optimized for FX10.

---

```

1:  $\hat{r}_c \leftarrow \{r_c, r_c\}$ 
2:  $\hat{c}_1 \leftarrow \{c_1, c_1\}$ 
3:  $\hat{u} \leftarrow \{1, 1\}$ 
4:  $\hat{dt} \leftarrow \{dt, dt\}$ 
5: for  $i = 1$  to  $N - 1$  do
6:    $\hat{\mathbf{q}}_i \leftarrow \{\mathbf{q}[i], \mathbf{q}[i]\}$ 
7:    $l \leftarrow \lfloor (\text{KeyPointer}[i + 1] - \text{KeyPointer}[i]) / 4 \rfloor$ 
8:    $k_0 \leftarrow \text{KeyPointer}[i]$ 
9:    $\hat{\mathbf{p}}_i \leftarrow \{\mathbf{0}, \mathbf{0}\}$ 
10:  for  $k = 0$  to  $l - 1$  do
11:     $k \leftarrow k_0 + 4k_1$ 
12:     $j_a \leftarrow \text{SortedList}[k]$ 
13:     $j_b \leftarrow \text{SortedList}[k + 1]$ 
14:     $j_c \leftarrow \text{SortedList}[k + 2]$ 
15:     $j_d \leftarrow \text{SortedList}[k + 3]$ 
16:     $\hat{\mathbf{q}}_A \leftarrow \{\mathbf{q}[j_a], \mathbf{q}[j_b]\}$ 
17:     $\hat{\mathbf{q}}_B \leftarrow \{\mathbf{q}[j_c], \mathbf{q}[j_d]\}$ 
18:     $\hat{\mathbf{r}}_A \leftarrow \hat{\mathbf{q}}_A - \hat{\mathbf{q}}_i$ 
19:     $\hat{\mathbf{r}}_B \leftarrow \hat{\mathbf{q}}_B - \hat{\mathbf{q}}_i$ 
20:     $\hat{r}_A^2 \leftarrow |\hat{\mathbf{q}}_1|^2$ 
21:     $\hat{r}_B^2 \leftarrow |\hat{\mathbf{q}}_2|^2$ 
22:     $\hat{w}_A \leftarrow \hat{u} - \lfloor \hat{r}_A^2 / \hat{r}_c^2 \rfloor$ 
23:     $\hat{w}_B \leftarrow \hat{u} - \lfloor \hat{r}_B^2 / \hat{r}_c^2 \rfloor$ 
24:     $\hat{f}_A \leftarrow [(24\hat{r}_A^6 - 48) / \hat{r}_A^{14} + 8\hat{c}_1] \hat{dt} \hat{w}_A$ 
25:     $\hat{f}_B \leftarrow [(24\hat{r}_B^6 - 48) / \hat{r}_B^{14} + 8\hat{c}_1] \hat{dt} \hat{w}_B$ 
26:     $\hat{\mathbf{p}}_i \leftarrow \hat{\mathbf{p}}_i + \hat{f}_A \hat{\mathbf{r}}_A + \hat{f}_B \hat{\mathbf{r}}_B$ 
27:  end for
28:   $\{\mathbf{p}_1, \mathbf{p}_2\} \leftarrow \hat{\mathbf{p}}_i$ 
29:   $\mathbf{p}[i] \leftarrow \mathbf{p}[i] + \mathbf{p}_1 + \mathbf{p}_2$ 
30:  for  $k = \text{KeyPointer}[i] + 4l$  to  $\text{KeyPointer}[i + 1] - 1$  do
31:     $j \leftarrow \text{SortedList}[k]$ 
32:     $\mathbf{r} \leftarrow \mathbf{q}[j] - \mathbf{q}[i]$ 
33:     $r^2 \leftarrow |\mathbf{r}|^2$ 
34:    if  $r^2 < r_c^2$  then
35:       $f \leftarrow [(24r^6 - 48) / r^{14} + 8c_2] dt$ 
36:       $\mathbf{p}[j] \leftarrow \mathbf{p}[j] + f \mathbf{r}$ 
37:    end if
38:  end for
39: end for

```

---

Nodes	Number of Particles	Elapsed Time [s]	Speed [MUPS]	Efficiency
1	8,000,000	351.821	22.7389	1.00
256	2,048,000,000	355.829	5755.57	0.99
512	4,096,000,000	362.775	11290.7	0.97
1024	8,192,000,000	368.275	22244.2	0.96
2048	16,384,000,000	364.272	44977.4	0.97
4096	32,768,000,000	373.007	87848.3	0.94
4800	38,400,000,000	364.920	105229.0	0.96

TABLE IV: Same as Tables II and III for the flat-MPI code optimized for FX10.

#### IV. MULTIBUBBLE NUCLEI

The study of multibubble nuclei is challenging, since it is a typical multiscale and multiphysics phenomenon that involves phase transitions at the microscopic scale and interbubble interactions at the meso- or macroscopic scale. From the viewpoint of computational simulations, phase transitions make simulations difficult since they involve the creation, movement, and annihilation of phase interfaces. If one simulates bubble nuclei using MD, then there is no need to consider phase transitions and phase interfaces. Both of these occur spontaneously in the simulation



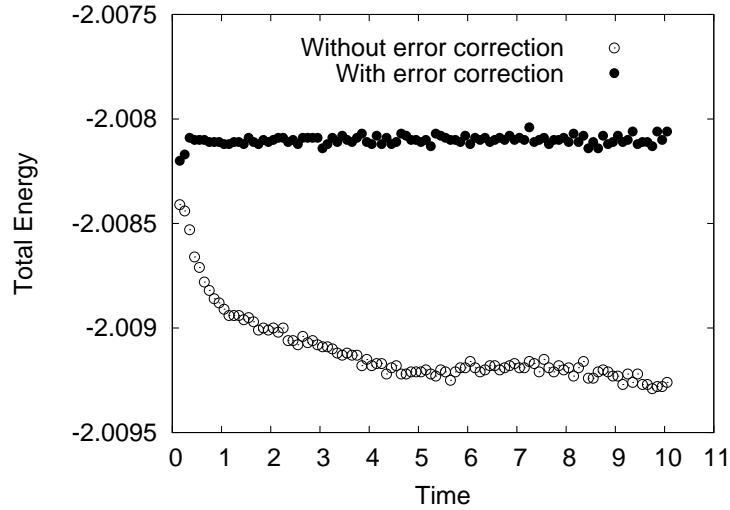


FIG. 4: Effect of error correction on reciprocal approximation. The number of particles is 4000. Computations are performed on a single core. The open circles denote the results without the error correction and the filled circles denote those with the error correction. The energy conservation with the error correction is sufficient.

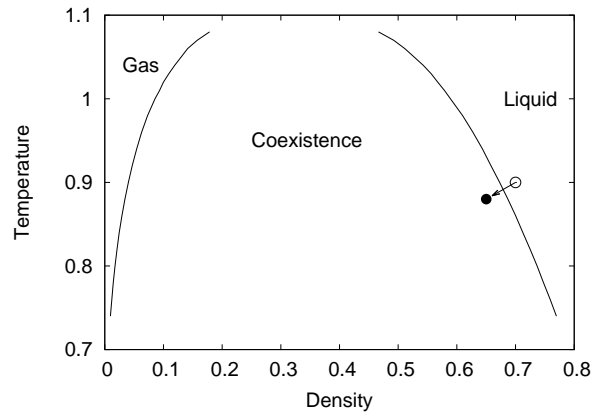


FIG. 5: Phase diagram of the system. The solid lines denote the gas-liquid coexistence densities. The open circle denotes the initial condition of the system. After the expansion, the state of the system moves to the point denoted by the solid circle.

as a result of the dynamics of particles. However, if one wants to reproduce multibubble nuclei by MD, then a huge number of particles are required; therefore, a large-scale simulation is necessary. There are several reports on extremely large-scale MD simulations. A team from Lawrence Livermore National Laboratory and IBM studied the Kelvin-Helmholtz instability by MD involving 62.5 billion particles on BlueGene/L, for which they won the 2007 Gordon Bell Prize [18]. In 2008, Germann and Kadau demonstrated one-trillion-atom MD on BlueGene/L using single precision [19]. However, to our best knowledge, there have been studies of large MD simulations involving a phase transition carried out on systems containing over 10,000 CPU cores.

In this study, we perform large-scale MD simulations of multibubble nuclei. There are two types of processes involving bubble nuclei; boiling and cavitation. Boiling occurs when a liquid is heated under a constant-pressure condition and cavitation occurs when the pressure of the liquid is reduced. We study cavitation phenomena since a simulation of cavitation requires only a thermostat, while that of boiling requires both a thermostat and a valostat.

Site	Scheme	Processes	Threads	System Size	Particles	Total Steps	Elapsed Time [h]
ISSP	Flat-MPI	1,024	1,024	$320 \times 320 \times 320$	22,937,600	690,000	10
ITC	Hybrid	4,800	76,800	$1,440 \times 1,200 \times 1,200$	1,449,776,020	178,000	15

TABLE V: Details of conditions. The parallelization scheme, the numbers of processes and threads used for the jobs, the dimensions of the system, the total number of particles, the total number of steps, and the elapsed time are shown for each site.

### A. Cavitation Process

We first maintain the system at a temperature  $T = 0.9$  with a density  $\rho = 0.7$  using a thermostat. This corresponds to the pure-liquid phase (see Fig. 5 for the phase diagram.). We use the Nosé-Hoover thermostat to control temperature [20]. Then we expand the system. The expansion is performed by changing the length scale of the system from  $l$  to  $\alpha l$ . The rescaling factor  $\alpha$  is chosen to be larger than unity. With this change, the position of each particle also changes from  $(q_x, q_y, q_z)$  to  $(\alpha q_x, \alpha q_y, \alpha q_z)$  while keeping the momenta. This procedure models a uniform and adiabatic expansion. We chose the rescaling factor to be 1.025. Then the density changes from 0.7 to 0.65, which is in the gas-liquid coexistence region in the phase diagram; therefore, bubbles start to appear. After expansion, we turn off the thermostat so that it does not affect the physical processes of the bubble nuclei. The integration scheme used for the isothermal time evolution is the second-order reversible system propagator algorithm (RESPA) [21], and the second order symplectic method is used for the microcanonical simulation with a time step  $\Delta t = 0.005$ . We perform a small job at ISSP and a large job at ITC. Hereinafter, we denote the smaller and larger jobs by ISSP and ITC, respectively. The conditions of the jobs are summarized in Table V.

### B. Bubble Identification

To identify bubbles, we divide the system into small subcells with a length of 3.0 and determine the local density for each subcell [22]. The densities of the gas and liquid coexisting in this system at  $T = 0.9$  are estimated to be 0.0402(2) and 0.6730(2), respectively [23]. Therefore, we define a subcell to be in the *gas state* when its density is less than 0.2. We have confirmed that the results do not change for other threshold values such as 0.3. We consider that neighboring *gas state* cells are in the same cluster and identify the bubbles using the site-percolation criterion on the simple cubic lattice. This process is called clustering, and the parallelization of clustering is quite expensive since it involves global communication similar to that of the fast fourier transform. Therefore, we gather all the data of local densities to the root node and perform clustering with serial calculation. To reduce memory consumption, we count the number of particles in each subcell and store it as an integer (data type `unsigned char` in C language) instead of storing the local density with floating-point variables. Since the volume of each subcell is 27 and the local density hardly exceeds 1.0, the number of particles in a subcell is at most 27, which can be stored in one byte without losing any information. A snapshot of the data is 1.4 MB at ISSP and 76.8 MB at ITC. To perform clustering, more work memory is required in addition to the data of the local densities.

### C. Simulation Results

We perform the job with flat-MPI at ISSP, whereas we perform the job with the hybrid at ITC in order not to use too much memory. The time evolution of pressure in each case is shown in Fig. 6. Although the pressure is positive when the system is in the pure liquid phase, it suddenly becomes negative after the expansion and increases as bubbles appear and grow. Snapshots obtained during the runs are shown in Figs. 7 and 8. In the job at ISSP, Ostwald-like ripening is observed after multibubble nuclei, and finally only one bubble remains in the system. Multibubble nuclei and Ostwald-like ripening are also observed at ITC.

The time evolutions of the total number of bubbles are shown in Fig. 9 (a). After the rapid increase in the number of bubbles, the number decreases slowly owing to interactions between bubbles, *i.e.*, small bubbles are eliminated owing to the growth of larger bubbles. The power-law decay of the number of bubbles can be observed. The distributions of the bubble sizes at the 10,000th step after expansion are shown in Fig. 9 (b). Although the results of both simulations show similar distributions, one can see that the distribution for the jobs at ISSP appears to be insufficient for further analysis.

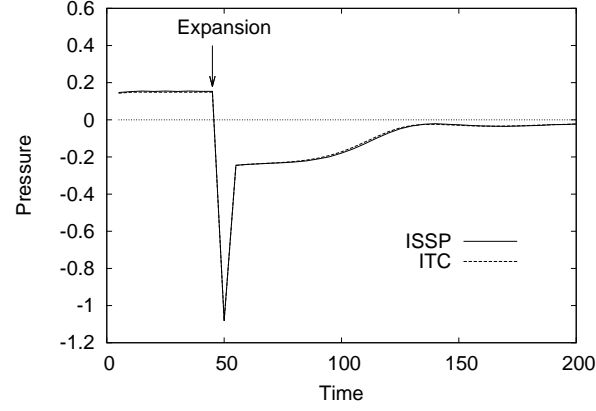


FIG. 6: Time evolutions of pressure. First the system is maintained in the pure liquid phase. At  $t = 50$ , the system expands and the pressure becomes negative. As bubbles appear and grow, the pressure increases after the expansion. The behaviors of the two jobs are almost identical.

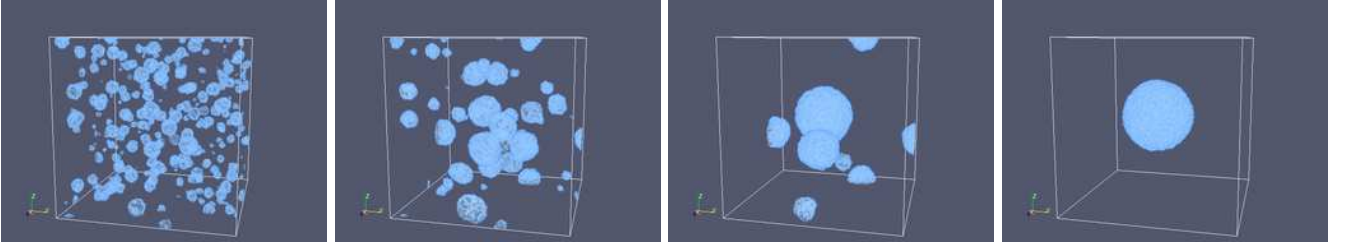


FIG. 7: (Color online) Snapshots of bubbles obtained during runs at ISSP. The direction of time evolution is from left to right. In the early stages, many bubbles appear and grow independently. Afterward, the bubbles start to interact and Ostwald-like ripening, where larger bubbles become larger and smaller bubbles disappear, is observed. Finally, one large bubble remains in the system. These pictures were produced by ParaView [24].

#### D. Load Imbalance

In the final part of this section, we discuss the load imbalance of our simulations. To check the validity of the pair list, we synchronize all threads at every time step. Therefore, the slowest thread determines the overall performance since the other threads must wait for the slowest thread due to the synchronization. The computational workload of a domain is roughly proportional to the number of particles it contains. Since the volumes of domains are identical, the workload is proportional to the number density. Since a liquid is almost incompressible, the largest density is determined by the coexisting density of the liquid at the temperature of the system. In our simulation, the temperature becomes 0.88 after the expansion. The densities of the coexisting gas  $\rho_g$  and liquid  $\rho_l$  at this temperature

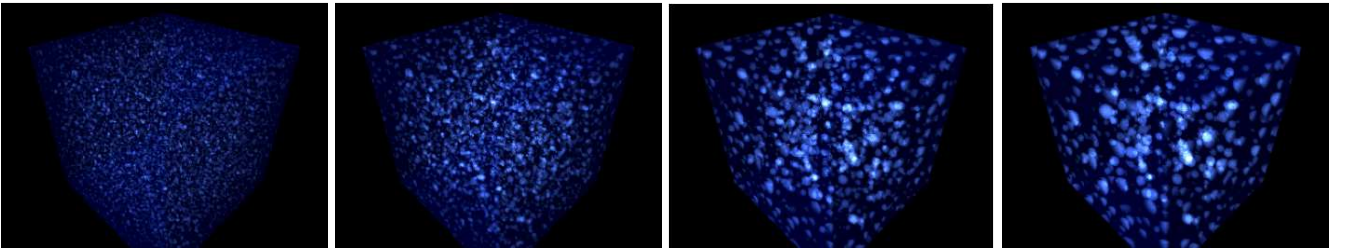


FIG. 8: (Color online) Snapshots of bubbles obtained during runs at ITC. The direction of time evolution is from left to right. Multibubble nuclei and the Ostwald-like ripening are observed. These pictures were produced by POV-Ray [25].

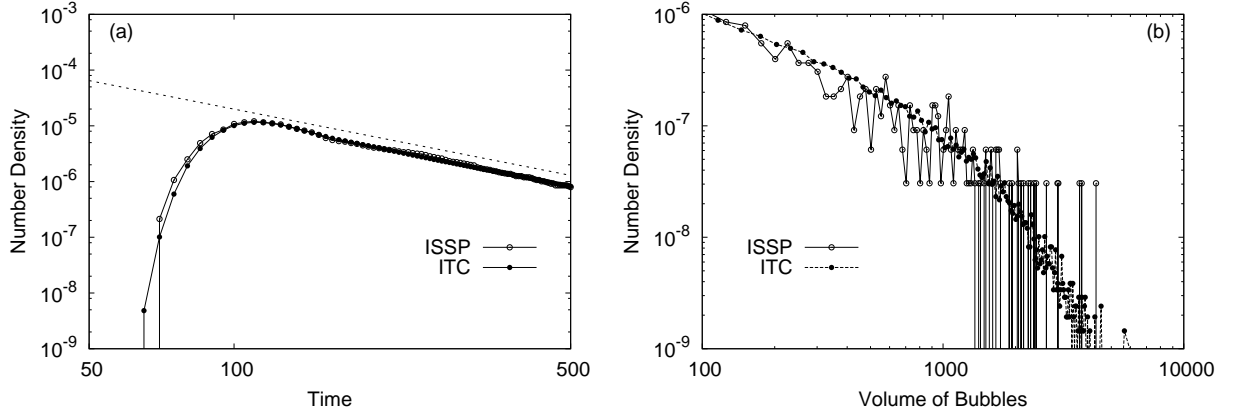


FIG. 9: (a) Time evolution of the total number of bubbles. The decimal logarithms are taken for both axes. The dashed line with a slope of -1.7 is a guide to the eyes. (b) Size distributions of bubbles at the 10,000th step after expansion. The decimal logarithms are taken for both axes.

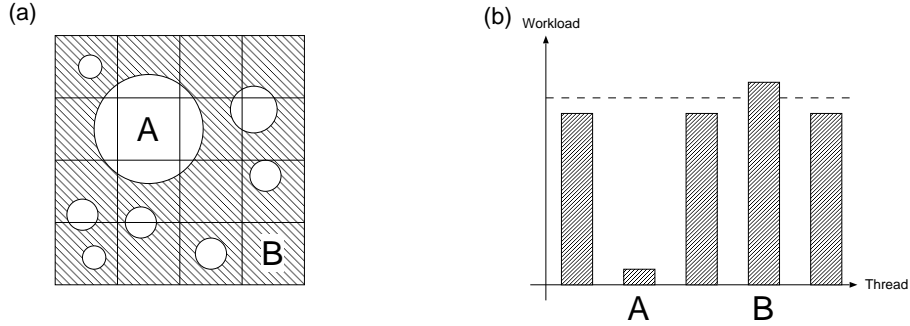


FIG. 10: Schematic illustrations of load imbalance. (a) Typical configuration of simulation of cavitation. There are many bubbles. The domain denoted by  $A$  is filled with pure gas, and the domain denoted by  $B$  is filled with pure liquid, *i.e.*, it does not contain any bubbles. (b) Workloads of threads are proportional to the densities of the assigned domains. The dashed line denotes the average density of the system, which is the workload of domains for a perfect load balance. The workload of thread  $B$  is about 20 times larger than that of thread  $A$ . However, it is only 5% larger than the average workload.

are  $\rho_g = 0.0344$  and  $\rho_l = 0.687$ , respectively [23]. Therefore, the computational time required for one step of the slowest process is proportional to  $\rho_l$ , which is 20 times larger than that of the fastest thread. If a perfect load balance is achieved, all threads have identical numbers of particles, and the densities in the domains become the same as the total number density  $\rho_{ave}$ . In our simulation, the average density after expansion is  $\rho_{ave} = 0.65$ . Therefore, the computational time of the simulation without load balancing is proportional to  $\rho_l$  and that with a perfect load balance is proportional to  $\rho_{ave} = 0.65$ . The reduction in speed due to the load imbalance is given by the ratio of the average workload to the heaviest workload, which is  $1 - \rho_{ave}/\rho_l \sim 0.05$ , *i.e.*, only 5% (see Fig. 10). This means that the load imbalance is minor problem in studies of cavitation phenomena. If one wants to study condensation phenomena, which involve droplet nuclei, then the problem of load imbalance becomes serious. Since the average density is close to that of the gas while the slowest thread is determined by the density of the liquid, the simulation performance can be 20 times worse than that of the perfectly load-balanced simulation.

## V. SUMMARY AND DISCUSSION

We have developed a pseudo-flat MPI MD code and performed benchmark and product runs utilizing up to 4800 nodes, and 76,800 cores. Benchmark simulations contained up to 38.4 billion particles and exhibited almost perfect scaling for flat-MPI. We studied cavitation phenomena with MD involving up to 1.45 billion particles and observed multibubble nuclei and Ostwald-like ripening. Bubbles appear as a result of the many-body interaction of particles,

and Ostwald-like ripening appears as a result of the many-body interaction of bubbles. This shows that direct simulations of multiscale physics at the atomic scale are now possible. In the case of more than 1 billion particles, the population dynamics of a bubble-size distribution can be observed with satisfactory statistics. This will allow the study of the nonequilibrium dynamics of multibubble interactions. We executed our codes on two different sites with different methods of parallelization, thus showing that our codes are portable and flexible. The flat-MPI method exhibited better performance than the hybrid, while the computation assigned to each core was identical. Another issue is to identify the main factor contributing to the parallel overhead of OpenMP. The amount of memory per computational power on a node will decrease as the number of cores per node increases. This trend will make it difficult to perform simulations with flat-MPI in the future. The pseudo-flat MPI approach allows us to perform a job with flat-MPI to obtain better performance or with the hybrid to reduce memory consumption. The developed code is available online [26].

### Acknowledgements

The authors would like to thank K. Nitadori, Y. Kanada, N. Kawashima, and S. Todo for valuable comments. This work was partially supported by Grants-in-Aid for Scientific Research (Contract No. 23740287) and by KAUST GRP (KUK-I1-005-04). The computational resource of Fujitsu FX10 was awarded by “Large-scale HPC Challenge” Project, Information Technology Center of the University of Tokyo. The computation was also carried out using the facilities of the Institute for Solid State Physics of the University of Tokyo.

- 
- [1] UPC language specifications are available from: <http://upc.gwu.edu>.
  - [2] The technical specifications of Co-Array Fortran are available from: [http://www.co-array.org/caf\\_def.htm](http://www.co-array.org/caf_def.htm).
  - [3] The official website of Titanium is <http://titanium.cs.berkeley.edu>.
  - [4] The latest version of the specifications of XcalableMP is available from: <http://www.xcalablemp.org>.
  - [5] H. Watanabe, M. Suzuki, and N. Ito, *Prog. Theor. Phys.* **126**, (2011) 203 – 235.
  - [6] S. D. Stoddard and J. Ford, *Phys. Rev. A* **8**, (1973) 1504 – 1512.
  - [7] M. P. Allen and D. J. Tildesley, *Computer Simulation of Liquids* (Clarendon Press, Oxford, 1987).
  - [8] D. E. Knuth, *The Art of Computer Programming* (Addison-Wesley, Reading, MA, 1973).
  - [9] G. S. Grest and B. Dünweg, *Comput. Phys. Commun.* **55** (1989) 269 – 285.
  - [10] D. M. Beazley and P. S. Lomdahl, *Parallel Comput.* **20** (1994) 173 – 195.
  - [11] B. Quentrec and C. Brot, *J. Comput. Phys.* **13** (1975) 430 – 432.
  - [12] R. W. Hockney and J. W. Eastwood, *Computer Simulation Using Particles* (McGraw-Hill, New York, 1981).
  - [13] L. Verlet, *Phys. Rev.* **159** (1967) 98 – 103.
  - [14] M. Isobe, *Int. Mod. Phys. C* **10** (1999) 1281–1293.
  - [15] S. Plimpton, *J. Comput. Phys.* **117** (1995) 1–19.
  - [16] M. J. Berger, M. J. Aftosmis, D. D. Marshall, S. M. Murman, *J. Parallel Distrib. Comput.* **65** (2005) 414 – 423.
  - [17] The official documents of MPI 2.2 are available from: <http://www.mpi-forum.org/docs/docs.html>.
  - [18] J. N. Glosli, D. F. Richards, K. J. Caspersen, R. E. Rudd, J. A. Gunnels, and F. H. Streitz, *Proceedings of the 2007 ACM/IEEE conference on Supercomputing* (2007) 58.
  - [19] T. C. Germann and K. Kadau, *Int. J. Mod. Phys. C* **19** (2008) 1315 – 1320.
  - [20] W. G. Hoover, *Phys. Rev. A* **31** (1985) 1695 – 1697.
  - [21] M. Tuckerman, B. J. Berne, and G. J. Martyna, *J. Chem. Phys.* **97** (1992) 1990 – 2001.
  - [22] H. Watanabe, M. Suzuki, and N. Ito, *Phys. Rev. E* **82** (2010) 051604.
  - [23] H. Watanabe, N. Ito, and C.-K. Hu *J. Chem. Phys.* **136** (2012) 204102.
  - [24] <http://www.paraview.org>
  - [25] <http://www.povray.org>
  - [26] The latest version of MDACP can be downloaded from: <http://mdacp.sourceforge.net>.